



# Cheryl's Hot Flashes #8

Cheryl Watson

Session 2543, August 23, 2002

Watson & Walker, Inc.

**publishers of 'Cheryl Watson's TUNING Letter'**

[www.watsonwalker.com](http://www.watsonwalker.com)

Technology ▪ Connections ▪ Results

# Cheryl's Hot Flashes

- User Experiences
  - WLM Velocities, Cache (HSB) performance, SORT & 64-bit, GRS
- Interesting Items
  - RedBooks, Hardware alternatives
- 6-Month Update
  - Z900 Cache, ERBSCAN, WLM Inits
- z/OS 1.4
  - Recommendation, My favorite and least favorite things
- This SHARE
  - NetView enclaves, Great sessions
- Soapbox Time
  - Short CPs, SHARE justification

# User Experiences

## ● WLM Velocities

- APAR **OW47277** for z/OS 1.1 changed how velocities are calculated (change was made for IRD). This code was then incorporated into later z/OS releases. Result: velocities are lower.
- So, **goals may need to be changed** when you migrate from any release of OS/390 to any release of z/OS.

# User Experiences

- WLM Velocities (cont.)

- Track APAR **OW55665** (OPEN, 7/18/02) which addresses the change.
- IBM intends to lower the impact of the velocity change and provide a utility to show you both velocities (track APAR and WLM Web site at [www.ibm.com/zseries/zos/wlm](http://www.ibm.com/zseries/zos/wlm)). So you may or may not need to change velocities.
- One question: how do you set velocities for policies that have mixed OS/390 and z/OS images?
- Reference: TUNING Letter, 2002, No. 4; & Cheryl's List #68: [www.watsonwalker.com/archives.html](http://www.watsonwalker.com/archives.html), session 2500 by **Kathy Walsh**

# User Experiences

## ● zSeries Cache Performance

### ■ G5/G6 CPU Cache

- Each CP has one 256K integrated Instruction and Data cache with a cache line size of 128 bytes
- When an instruction modifies data within 128 bytes, there is some associated overhead

### ■ zSeries CPU Cache

- Each CP has two 256K caches, one for Instruction (read-only) and one for Data, each with a cache line size of 256 bytes
- When an instruction in the I-cache modifies data within 256 bytes, there is significantly more overhead

# User Experiences

## ● Cache Performance – z900

- The z900 design was changed to separate the Data cache and Instruction cache for each CP. The cache line size is 256 bytes.
- If you modify something within a line that has already been pulled into the instruction cache, it must be moved back over to the data cache. The invalidation takes some overhead.
- Some code was identified (examples are SAS, an ADABAS utility, some IBM utilities) that ran longer on the z900 because of this behavior. So ISVs have provided performance fixes to resolve this.
- Reference: Cheryl's List #50 & #51 at [www.watsonwalker.com/archives.html](http://www.watsonwalker.com/archives.html)

# User Experiences

## ● Cache Performance – G5/G6

- IBM G5/G6 processors use a 128-byte cache entry, but processing is slightly different.
- **Lawrence Jermyn** from Fidelity Investments and **Dave Brown** of IBM/Dallas made a really neat discovery: modifying data outside 256 bytes on a G5/G6 \*also\* provides performance improvements (although not as much) as shown in the following test results:

	Raw CPU Seconds		
	Close	Dispersed	
<b>G6</b>	<b>21</b>	<b>4.7</b>	<b>78% drop</b>
<b>z900</b>	<b>103.7</b>	<b>3.2</b>	<b>97% drop</b>

- Some ISVs have confirmed this from their own tests.
- Reference: TUNING Letter, 2002, No. 4

# User Experiences

## ● Cache Performance – MP3000

- **Peter I. Vander Woude** from Harris Teeter ran the same test on a Multiprise 3000/H70 (which is based on a G5 chip). His results:

- **Raw CPU seconds:**

**Close - 32.44; Dispersed - 7.30 = 77% drop**

## ● Cache Performance – Conclusion

- Fixes for z900 cache problem are good for G5/G6 and MP3000 machines. Not as great as on zSeries, but worth the effort. Get fixes from software vendors (IBM, SAS, Software AG, CA, etc.)



# User Experiences

## ● Sort and 64-Bit

### ■ Flash10185 - z/Architecture Migration Checklist by Bette Brody

- "DFSORT considerations - Possible system performance impact (paging) when using in 64-bit mode. Check DFSORT maintenance using the search D/T2064. The **DFSORT Installation and Customization Release 14 publication (SC33-4034)** describes the hiperspace installation options referenced below.
- DFSORT users who specify any one of the following ICEMAC (installation) options: **HIPRMAX=OPTIMAL EXPMAX=MAX EXPOLD=MAX EXPRES=MAX HIPRMAX=p% EXPMAX=p% EXPOLD=p% EXPRES=p%** may see an increased usage of hiperspace on the z/900 CPU. This is due to calculations now being based on central storage values instead of expanded storage values. **Increased paging may occur.**"

# User Experiences

## ● Sort and 64-Bit

- One site had an aux storage shortage when SYNCSORT with CORE=MAX caused other address spaces to page out heavily.
- Another site got AFC (available frame count) under 10 and 5-30 minute TSO logons when four DFSORTs with HIPERMAX=OPTIMAL were running.
- The DFSORT development team recommends using HIPERMAX=OPTIMAL unless you run into paging problems.
- If you think there's a problem, contact your sort vendor.

# User Experiences

## ● GRS on z/OS 1.2

- One site reported a large CPU increase when going to z/OS 1.2 (CPU busy from 70% to 90%). Problem was found to be in GRS exit ISGNQXIT. IBM gave user fixes that worked, but no APAR was opened, since the fixes were going to be incorporated in another release of GRS.
  - The change probably came from APAR **OW49779**, which added wildcard names to GRS. (My opinion only, but if true, it could also impact OS/390 R8/R10 users who apply the APAR.)
  - The wildcard support was added to eliminate most of the exits.
- Another interesting APAR is **II13047**: FIXES TO ADDRESS A HOLE CLOSED BY GRS WILDCARD SUPPORT
- Look for pre-reqs for CA-MIM as well

# Interesting Items

- Neat New Redbook

- **SG24-6542-00 - z/OS for e-business: An Introduction to System Tuning, 8/2/02 – great for WebSphere tuning**

- Investigate Hardware Alternatives

- FLEX-ES provides z/OS and OS/390 emulation software for PCs and laptops. They're the second largest mainframe vendor! See [www.funsoft.com](http://www.funsoft.com) for details and list of hardware vendors. You can have an MP3000 that runs 64-bit! Some large companies are investigating these options as sysprog sandboxes.
- z/OS.e can be used to run new apps (WebSphere, Domino, etc.) at a fraction of the cost. Look into the possibilities.

# 6-Month Update

---

- Based on my Hot Flashes #7

- See session 2543 from Nashville Proceedings – [www.share.org](http://www.share.org)
- Here are some updates
- Note: Nashville proceedings move to the Members-Only section as soon as the San Francisco proceedings are online

# 6-Month Update

## ● z900 CP High Speed Cache

- **Last SHARE:** When ADABAS moved to z900 from a G6, the CPU time increased significantly (should have decreased). There is probably some code in ADABAS that touches data within 256 bytes of instruction that hasn't been identified or fixed yet.
- **Update:** The situation was caused by a single ADABAS utility, ADACMP. Software AG provided a fix. You'll find additional 64-bit updates at [www.softwareag.com/corporat/prodstat/ibm/default.htm](http://www.softwareag.com/corporat/prodstat/ibm/default.htm). Thanks to Software AG for one of the easiest references for architecture changes. These updates could also be good for G5, G6, and MP 3000 machines.

# 6-Month Update

## ● ERBSCAN

- **Last SHARE:** ERBSCAN introduced with RMF OS/390 R4. Allows you to browse SMF records via TSO, with some formatting for RMF records. Few people know about it, but it's VERY handy!
- **Update:** Not documented until OS/390 R10: **RMF Programmer's Guide - SC28-1951 (OS/390 R10) and SC33-7994 (z/OS)**. One item not documented: ERBSCAN can only go against a sequential (offloaded) data set. If you go against the SMF VSAM file, nonsensical errors occur. RMF APAR **OW55788** will add an error message if you try to access SMF VSAM file (OPEN, target 9/30/02). Thanks to **Tom Conley** of Pinnacle Consulting Group.

# 6-Month Update

## ● WLM Inits

- **Last SHARE:** One site recently had to move back to compat mode because the INITs were getting dispatched at too high a priority on a busy system and the more important onlines were being impacted. WLM doesn't allow the user to define or reset \*MASTER\*, INIT, or WLM.
- **Update:** Inits run at PVLDP in compat mode and SYSSTC in goal mode. The INIT can use a large amount CPU, especially when there are a lot of DDs: the ACS routines are run here for allocations, SMF user exits are run here, and a SAF call for account code verification is also run here. WLM is investigating a solution (was discussed at EWCP sessions 2505 and 2518). No APAR assigned yet, but I'll let you know.



# z/OS 1.4

---

- My previous recommendation:

Stay on OS/390 until you go to a zSeries machine

# z/OS 1.4

- My new recommendation:

## Migrate to z/OS 1.4 now!

- Pricing was changed to allow you the choice of WLC or PSLC (take your time to determine which method works best for you, but realize that it may be worth the effort to investigate it)
- Bimodal support will be added 9/27/02 (this removes the risk of switching out machines and failing due to missing fixes)
- Use migration guide for OS/390 R10 to z/OS 1.4 (follow the crowd)
- Be sure to migrate to goal mode in R10 well before going to z/OS 1.4 (to get used to it)
- Check out z/OS FAQs at [www.ibm.com/zseries/faq/](http://www.ibm.com/zseries/faq/)

# z/OS – My Favorite Things

- z/OS 1.2 SDSF Enclave Support
  - SDSF Enclave Display
    - New ENC display screen to show enclaves
  - SDSF Enclave Service Class Reset
    - SDSF SPE (**PQ50025**) adds the subsystem to the ENC display and allows an operator to reset an enclave to a different service class, to quiesce an enclave, or to resume an enclave by overtyping on the ENC display
  - Also mentioned in sessions 2500 (**Kathy Walsh**) and 2515 (**Jürgen Holtz**).

# z/OS – My Least Favorite Things

- z/OS 1.2 SDSF Enclave Support (cont.)
  - But I'd like SDSF to make a change.
  - To understand why I say this, you need to know that there are independent enclaves and dependent enclaves.
    - Independent enclaves are classified by WLM and assigned to a service class (e.g DDF, OMVS).
    - Dependent enclaves run in the service class of their originator (e.g. TSO making an SQL call or CICS transactions issuing a Stored Procedure to DB2).

# z/OS – My Least Favorite Things

- z/OS 1.2 SDSF Support (cont.)
  - WLM provides an API to allow a program to reset an enclave into a new service class (or quiesce/resume it), but does not allow the reset to be issued on a dependent enclave.
  - SDSF, however, lets the user reset dependent enclaves.
  - This could be disastrous if they don't know what they're doing!
  - A better solution would be to display a warning that you're about to change a dependent enclave and indicate which work will be changed if you confirm the reset.
  - It looks like a requirement to me!

# z/OS – My Favorite Things

- z/OS 1.3

- ASM and PAVs

- ASM utilizes dynamic PAV aliases if the page data set is on an ESS. (Prior to this, paging volumes weren't eligible for dynamic PAVs). A change in ASM lets it make multiple requests to the same paging device via PAVs, so that one PAV can be reserved for page faults and the other PAV can do both reads and writes. So an SVC dump shouldn't stop page-in requests for other work.

- No support for compat mode

- Important recommendation: get to goal mode BEFORE going to R3 or R4! IBM provides an Install Definition Utility, IWMINSTL, that can install a service policy in batch. They also provide a sample service definition in SYS1.SAMPLIB(IWMSSDEF), but it requires modifications.

# z/OS – My Favorite Things

## ● z/OS 1.3

### ■ Report Class Period Enhancement

- If only one service class is assigned to a report class, then the report class can collect all of the period, goal, and sample data that the service class does. This allows you to have a multi-period TSO service class and still get reporting for different groups of users. **Wow!**
- Trick: if you want to get CICS/IMS transaction distributions, then assign them to a service class and report class, but set the region to REGION managed. You'll need to define a default service class as well. There will be nothing running in the service classes, but this allows the report class to pick up the goals of the unused service classes for goal calculation. Thanks to **Jürgen Holtz** (IBM) for this suggestion.

# z/OS – My Favorite Things

## ● z/OS 1.3 - New SMF Records

- **Barry Merrill** noted this in his CHANGES documentation. Thanks, Barry! z/OS 1.3 added two significant records to the SMF file that will prove to be invaluable! But IBM forgot to show them as a change to SMF in the z/OS 1.3 manuals. Silly them!
- New SMF type 42, subtype 20, is written when STOW INITIALIZE is used to delete all members of a PDSE. The record shows the data set that was initialized, along with the job and user doing the initialization.



# z/OS – My Favorite Things

## ● z/OS 1.3 - SMF Records (cont.)

- More importantly, a new SMF type 42, subtype 21, is written when any member of a PDS or PDSE is deleted! The record shows the job and user performing the deletion, along with the member deleted, its data set name, and any aliases that were also deleted.
- MXG users can find this new information in datasets TYPE4220 and TYPE4221 starting with MXG 20.02. You can also find the subtype descriptions in macro IGWSMF in SYS1.MACLIB.
- Reference: TUNING Letter, 2002, No. 4

# z/OS - More Least Favorite Things

---

- z/OS 1.4 - Duplicate TSO Logons

- JES2 no longer checks TSO logons for duplicates (from session 2800, **Riaz Ahmed**, Hot Topics).
- You can turn it off, but 1.4 is shipped allowing duplicates!
- Why did they do this?

# This SHARE

- NetView V1.4 Can Use Enclaves

- But don't follow the NetView manual for setting goals; instead use APAR **OW54858** (not on database yet).
- If you turn on enclaves for NetView (by removing \* from \*WLM.SubSystemName=&DOMAIN in DSIPARM member CNMSTYLE), and don't classify all transactions, they'll run in SYSOTHER. (I reported this in my TUNING Letter, 2002, No. 3; **Kathy Walsh** described it in session 2500.)
- But what isn't documented: the valid classification rules for subsystem NETV are LU, PRI, SI, TC, TN, UI. Thanks to **Jürgen Holtz** for this list. If you use a different rule (e.g. AI), the work will also be assigned to SYSOTHER (because WLM doesn't check the rules for NETV).

# This SHARE

## ● Great Session Materials

- Session 2588 on Tuning DFSORT by **Carrie Van Noorden**, IBM. [www.ibm.com/storage/dfsort](http://www.ibm.com/storage/dfsort) (look for DFSORT papers)
- Session 2500 on z/OS and OS/390 Performance "Hot" Topics by **Kathy Walsh**, IBM. This is always superb!
- Session 2519 on Migrating to Fewer/Faster CPs in an LPAR Environment by **Walt Caprice**, IBM.
- Session 2809 on CF Structure Duplexing by **Mark Brooks**, IBM. See CF duplexing white paper at [www.ibm.com/zseries/pso](http://www.ibm.com/zseries/pso).
- Session 2810 on zSeries Parallel Sysplex Update by **Angelo Corridori**, IBM. Interesting fee offering: Parallel Sysplex Trainer (provides a fully functional parallel sysplex in as little as 3 hours); includes training. Easy way to get into parallel sysplex.

# This SHARE

- Great Session Materials

- Session 3061 on VSAM Performance Tuning by **Ron Ferguson**, Mainstar Software.
- Session 2571 'WLM – After Goal Mode, the Yellow Brick Road Didn't Lead to the Emerald City' by **Norman Hollander**, Candle. He includes an extensive list of WLM APARs at the end.
- Session 2818, 'Zee Guide to z/Architecture', by **Bob Shannon**, John Hancock. How to program in 64-bit.

# Soapbox Time

- Short CPs

- See session 2510 by **Walt Caprice**.
- His session should be required attendance for every installation.
- Two problems:
  - Some sites assign too many LPs to an LPAR (that's about 70% of the installations that I see)!
  - Sites only plan total MIPS capacity for new machines and forget about LP capacity.
- For existing systems, reducing the number of LPs will reduce LPAR overhead AND increase the amount of MIPS available to an LP.

# Soapbox Time

- Short CPs

- For upgrade planning, you must ensure three things:
  - The total number of MIPS is equal to or greater than the current configuration
  - The MIPS per LP is equal to or greater than the current configuration.
  - You've accounted for additional LPAR overhead if the number of LPs to CPs exceeds 3 to 1.
- If you don't do these, you could run into severe performance problems!

# Soapbox Time

## ● Short CPs

### ■ Example on z900-1C8:

- LPAR1 has 8 LPs with weight of 75
- LPAR2 has 8 LPs with weight of 25
- Z900-1C8 has 1620 MIPS (203 MIPS per CP);
  - LPAR1 gets 1215 MIPS ( $1620 * .75$ )
  - LPAR2 gets 405 MIPS ( $1620 * .25$ )
- MIPS per LP:
  - LPAR1 gets  $1215/8 = 152$  MIPS per LP
  - LPAR2 gets  $405/8 = 51$  MIPS per LP



# Soapbox Time

- Short CPs

- Example (cont.)

- BUT: LPAR1 only needs 7 LPs and LPAR2 only needs 3 LPs.
    - If you reduce the LPs, you reduce the LPAR overhead and each LPAR gets more CPU every time it's dispatched:
      - $\text{LPAR1} = 1215/7 = 174 \text{ MIPS per LP}$
      - $\text{LPAR2} = 405/3 = 135 \text{ MIPS per LP}$
    - This should be a no-brainer! Even IRD can do it better!
    - Do these calculations **BEFORE** you order a new machine!

# Soapbox Time

- You Need to Justify SHARE
  - Travel budgets have been slashed
  - Managers think that conferences are a boondoggle (because many of the ones THEY go to **ARE**)
  - So do BIG SHARE trip reports (I used to create a 100-page 'book' and route it to the world) to show them the value
  - Implement at least one change you've heard of and document the savings
  - Compute the cost of a high-priced consultant and show them how much free hallway consulting you've received
  - Compare to many of the education courses and you'll see that SHARE is quite a bargain!

# Questions

---



- Email: [cheryl@watsonwalker.com](mailto:cheryl@watsonwalker.com)
- Web site: <http://www.watsonwalker.com>